

Gittins Index Theorem

Thomas Kesselheim

Last Update: June 15, 2020

Suppose you are the boss of a large company. Your company is producing a well-running product out of which you make a profit of 10 units of money per year. You can keep the production running as it is, or you can invest the profits into research and development. Chances are that after five years of development, you can decrease the production cost and make a profit of 11 units of money per year. However, you have to spend your entire profits during this time and there is still a 50 % chance that after five years you realize that the development failed. Then, of course, you can still keep your production running, keeping the profit of 10 units of money per year. Under what circumstances would you decide to invest?

Such a problem can be modeled as a multi-armed bandit, which we introduced last time. Today, we will show what their optimal policies look like. The theory of Markov decision processes applies and, as we have an infinite time horizon, one could compute an optimal policy via linear programming, value iteration, and so on. However, the state space will be quite huge and the optimal policies have a particularly nice structure.

1 Recap: Markovian Multi-Armed Bandits

To define a multi-armed bandit, we first define a *single-armed bandit*. This is a Markov decision process that has only two actions $\mathcal{A} = \{\text{play}, \text{pause}\}$. The state transitions and rewards for action **play** are arbitrary, but $p_{\text{pause}}(s, s) = 1$, $r_{\text{pause}}(s) = 0$ for all $s \in \mathcal{S}$. That is, when using action **pause**, the process remains in its state and gives no reward.

A *multi-armed bandit* is a parallel composition of such single-armed bandits. We have $\mathcal{S} = \mathcal{S}_1 \times \dots \times \mathcal{S}_n$, where \mathcal{S}_i is the state space of the i^{th} single-armed bandits. Available actions are $\mathcal{A} = \{\text{play}_1, \dots, \text{play}_n, \text{pause}\}$, where play_i means that we run the **play** action on the i^{th} single-armed bandit and **pause** on any other. So the different single-armed bandits operate independently but we may only play one arm at a time.

Note that when a Markovian policy (for example an optimal policy) decides to pause, it remains in the state and therefore will keep pausing from now on and not resume playing an arm again. If $\gamma = 1$ then it would be irrelevant in which order we play the arms. However, because $\gamma < 1$, time is the distinguishing factor.

We consider the infinite time-horizon setting with discounts, so for a policy π

$$V(\pi, s_0) = \mathbf{E} \left[\sum_{t=0}^{\infty} \gamma^t r_{a_t^\pi}(s_t^\pi) \right] .$$

Recall that the value of an optimal policy is given as

$$V^*(s) = \max_{a \in \mathcal{A}} \left(r_a(s) + \gamma \sum_{s' \in \mathcal{S}} p_a(s, s') V^*(s') \right) .$$

2 Optimal Policy

In order to describe the optimal policy, consider the following problem with a single arm i . Suppose you can only choose between playing this arm or pausing and you had to pay λ every

time you played the arm. Then the maximum expected reward that you could get from arm i with an optimal policy is

$$V_i(s, \lambda) = \max \left\{ 0, r_{\text{play},i}(s) - \lambda + \gamma \sum_{s' \in \mathcal{S}} p_{\text{play},i}(s, s') V_i(s', \lambda) \right\} .$$

Note that $V_i(s, \lambda)$ only depends on the state of arm i , not on the states of the other arms.

Observe that for larger charges λ the value $V_i(s, \lambda)$ gets smaller and smaller. This means, there is some amount $\delta(s)$ that makes the optimal policy only exactly as good as not playing at all. Formally,

$$\delta_i(s) = \sup\{\lambda \mid V_i(s, \lambda) > 0\} = \inf\{\lambda \mid V_i(s, \lambda) = 0\} .$$

This is the *fair charge* or the *Gittins index* of arm i in state s .¹

Example 14.1. In Figure 1, we see one arm with deterministic transitions. In state E, the fair charge is 0. In states C and D, it is 1: We play once (or maybe twice from C) and then stop. For state B, we play once if the charge is 10. Finally, state A is interesting. The fair charge is 4: The reward from playing twice with charge λ is $1 - \lambda + \frac{1}{2}(10 - \lambda)$, which is 0 exactly for $\lambda = 4$.




Figure 1: A simple example of one deterministic arm with fair charges for $\gamma = \frac{1}{2}$.

Our main result for today is the following theorem.

Theorem 14.2. It is an optimal policy to always play the arm with the highest Gittins index.

3 Single-Armed Bandit with Charges

We first consider the problem with only a single arm. Based on the fair charges $\delta(s)$, it is very easy to describe an optimal policy for the arm with charge λ : Whenever in a state s with $\delta(s) \geq \lambda$ choose **play**, whenever in a state s with $\delta(s) < \lambda$ choose **pause**.²

We can also bound the reward of any policy by the fair charges of the states that it enters during its execution.

Lemma 14.3. Consider a policy for a single arm that first only chooses **play** and then only chooses **pause**. Let τ be the index of the step in which it chooses **pause** for the first time. Then

$$\mathbf{E} \left[\sum_{t=0}^{\tau-1} \gamma^t r_{\text{play}}(s_t) \right] \leq \mathbf{E} \left[\sum_{t=0}^{\tau-1} \gamma^t \min_{t' \leq t} \delta(s_{t'}) \right]$$

with equality if $\delta(s_\tau) = \min_{t' \leq \tau} \delta(s_{t'})$ with probability 1.

Proof. Let us first consider the case of a policy for which $\delta(s_\tau) = \min_{t' \leq \tau} \delta(s_{t'})$ with probability 1. So, this is a policy which only stops playing when reaching a state that has the smallest fair charge it has seen so far. Any execution of such a policy naturally decomposes into phases of




Figure 2: Phases in the proof of Lemma 14.3.

random length: Let $\tau_0 = 0$ and let τ_{k+1} be the first time step $t > \tau_k$ at which $\delta(s_t) < \delta(s_{\tau_k})$ (see Figure 2). The policy chooses **pause** at some point in time τ_k .

An alternative way to view the phases is as follows. Fix everything that happens until time τ_k . At time τ_k , start an optimal policy for the arm with charges $\delta(\tau_k)$. This optimal policy stops exactly at time τ_{k+1} because this is the first time that the charge exceeds the fair charge. The expected reward of this optimal policy starting at time τ_k with charge $\delta(s_{\tau_k})$ is exactly 0. So,

$$\mathbf{E} \left[\sum_{t=\tau_k}^{\tau_{k+1}-1} \gamma^{t-\tau_k} (r_{\text{play}}(s_t) - \delta(s_{\tau_k})) \mid \tau_k \right] = 0 .$$

Equivalently,

$$\mathbf{E} \left[\sum_{t=\tau_k}^{\tau_{k+1}-1} \gamma^t r_{\text{play}}(s_t) \mid \tau_k \right] = \mathbf{E} \left[\sum_{t=\tau_k}^{\tau_{k+1}-1} \gamma^t \delta(s_{\tau_k}) \mid \tau_k \right] = \mathbf{E} \left[\sum_{t=\tau_k}^{\tau_{k+1}-1} \gamma^t \min_{t' \leq t} \delta(s_{t'}) \mid \tau_k \right] .$$

Taking the sum over all k , the equality in the lemma follows.

To get the upper bound for a general policy, we can follow the argument above with the exception that the policy might stop one of its sub-phases early. In this case, the fair cost of the current state is higher than the charge. This means that expected reward for this sub-phase is at most 0. \square

Lemma 14.4. *Consider an arbitrary policy for a single arm and let the indices of steps in which it plays the arm be denoted by T (possibly random, depending on previous states). Then*

$$\mathbf{E} \left[\sum_{t \in T} \gamma^t r_{\text{play}}(s_t) \right] \leq \mathbf{E} \left[\sum_{t \in T} \gamma^t \min_{t' \leq t} \delta(s_{t'}) \right]$$

with equality if $\delta(s_t) = \min_{t' \leq t} \delta(s_{t'})$ for all $t \notin T$ with probability 1.

Proof. Note that Lemma 14.3 is exactly the case that $T = \{0, 1, \dots, \tau - 1\}$.

It is easy to extend the lemma to the case $T = \{t', \dots, t' + \tau - 1\}$ because then $\delta(s_0) = \dots = \delta(s_{t'})$ and both sides get multiplied by the same $\gamma^{t'}$.

In general, T can be considered a union of disjoint intervals, each of which has the form $\{t', \dots, t' + \tau - 1\}$. By adding up the resulting inequalities, the lemma follows. \square

¹The original definition by Gittins and Jones is a little different but has the same consequences.

²For $\delta(s) = \lambda$ actually both choices are equally good.

4 Proof of Gittins Index Theorem

Proof of Theorem 14.2. To prove the theorem, let T_i be the set of steps in which the Gittins index policy plays arm i . Let us observe how the Gittins index $\delta_i(s_t)$ changes over time. If $t \notin T_i$, then $\delta_i(s_{t+1}) = \delta_i(s_t)$. If $t \in T_i$ then $\delta_i(s_{t+1})$ can differ from $\delta_i(s_t)$. If it gets larger, then we keep playing the arm. We only stop playing the arm when its index falls below the value that we started from, meaning it is an all-time low. In other words, if $t \notin T_i$ then $\delta_i(s_t) \leq \min_{t' \leq t} \delta_i(s_{t'})$.

This allows us to invoke Lemma 14.4. We know that the expected reward from playing arm i is exactly

$$\mathbf{E} \left[\sum_{t \in T_i} \gamma^t \min_{t' \leq t} \delta_i(s_{t'}) \right]$$

and so the overall expected reward is exactly

$$Q = \sum_{i=1}^n \mathbf{E} \left[\sum_{t \in T_i} \gamma^t \min_{t' \leq t} \delta_i(s_{t'}) \right]$$

For any other policy, we get different values of T_i and $\delta_i(s_t)$ but by Lemma 14.4 its expected reward is still upper-bounded by the respective Q . Therefore, it is sufficient to show the following proposition.

Proposition 14.5. *Among all policies, the Gittins index policy maximizes*

$$Q = \mathbf{E} \left[\sum_{i=1}^n \sum_{t \in T_i} \gamma^t \min_{t' \leq t} \delta_i(s_{t'}) \right] .$$

We compare an arbitrary policy π to the Gittins index policy. For simplicity of the argument, we assume that both policies play each arm infinitely often. The spirit of the argument does not change without this assumption but things get much more messy.

For policy π as well as for the Gittins index policy, let us denote by x_t or y_t respectively, the value of $\min_{t' \leq t} \delta_i(s_{t'})$ for the arm i chosen in step t .

Arm i randomly transitions from one state in \mathcal{S}_i to another one, every time it is played. Let us fix these transitions arbitrarily. This way, the sequences x_0, x_1, \dots and y_0, y_1, \dots are not random anymore but fixed. Furthermore, because we assume that each arm is played infinitely often, they contain exactly the same numbers because each arm makes the same state transitions, only the order varies.

For the Gittins index policy, we have $y_0 \geq y_1 \geq \dots$, so the sequence is non-increasing. Therefore, now

$$Q^\pi = \sum_{t=0}^{\infty} \gamma^t x_t \leq \sum_{t=0}^{\infty} \gamma^t y_t = Q^{\text{Gittins}} .$$

This holds for any fixed transitioning of each single arm, so it also holds in expectation. \square

Further Reading

- On the Gittins Index for Multiarmed Bandits, R. Weber, Ann. Appl. Probab. (This proof without formulas)
- Four proofs of Gittins' multiarmed bandit theorem, E. Frostig, G. Weiss, Applied Probability Trust (This and other proofs, with heavy notation)